Strategic Budget Selection in a Competitive Autobidding World

Yiding Feng
yiding.feng@chicagobooth.edu
University of Chicago
Chicago, IL, USA

Brendan Lucier
brlucier@microsoft.com
Microsoft Research
Cambridge, MA, USA

Aleksandrs Slivkins
slivkins@microsoft.com
Microsoft Research
New York City, NY, USA

ABSTRACT

We study a game played between advertisers in an online ad platform. The platform sells ad impressions by first-price auction and provides autobidding algorithms that optimize bids on each advertiser's behalf, subject to advertiser constraints such as budgets. Crucially, these constraints are *strategically* chosen by the advertisers, and define an "inner" budget-pacing game for the autobidders. Advertiser payoffs in the constraint-choosing "metagame" are determined by the equilibrium reached by the autobidders.

Advertiser preferences can be more general than what is implied by their constraints: we assume only that they have weakly decreasing marginal value for clicks and weakly increasing marginal disutility for spending money. Nevertheless, we show that at any pure Nash equilibrium of the metagame, the resulting allocation obtains at least half of the liquid welfare of any allocation and this bound is tight. We also obtain a 4-approximation for any mixed Nash equilibrium or Bayes-Nash equilibria. These results rely on the power to declare budgets: if advertisers can specify only a (linear) value per click or an ROI target but not a budget constraint, the approximation factor at equilibrium can be as bad as linear in the number of advertisers.

CCS CONCEPTS

• **Theory of computation** → **Algorithmic game theory**.

KEYWORDS

sequential auction, metagame, liquid welfare, price of anarchy

ACM Reference Format:
Yiding Feng, Brendan Lucier, and Aleksandrs Slivkins. 2024. Strategic Budget Selection in a Competitive Autobidding World. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC '24), June 24–28, 2024, Vancouver, BC, Canada.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3618260.3649688

1 INTRODUCTION

In many large online platforms, it is increasingly common for users to delegate their choices to algorithmic proxies that optimize on their behalf. Examples include autobidders in online advertising markets [30, 41, 42], dynamic price adjustment algorithms for sales or rental platforms like Amazon and Airbnb [2, 4], price prediction tools for flights [26, 32], and more. These tools typically employ techniques from machine learning to optimize some goal subject to user-specified constraints or targets. When many users deploy these algorithmic tools simultaneously, the algorithms are effectively competing against each other. But in addition to the game between the algorithms, this setup defines a metagame between the users, who are choosing the parameters of their algorithms in anticipation of the competition. This raises a natural question: *how will users play the metagame, and what is the impact on the market?*

We study these questions through the lens of autobidding in online advertising markets. Online advertising platforms sell individual advertising events (like clicks or conversions) by auction, with a separate auction held for each ad impression. These auctions are strategically linked and optimal bidding is complicated, but advertisers can delegate the bidding details to an autobidder. We focus on online advertising in part because the use of automated bidding algorithms is very well established in that market, with all major advertising platforms offering integrated autobidding services [30, 41, 42].

The autobidding paradigm involves three layers: advertisers, autobidders, and auctions. See Figure 1 for an illustration. Each advertiser has their own individual autobidding algorithm (often provided by the platform) that bids on their behalf. An advertiser provides their autobidder with aggregate instructions like a budget constraint and/or a maximum bid.[1] The autobidder then participates in many individual auctions, implementing learning algorithms that use auction feedback to adjust bids in order to maximize the total value received (e.g., total clicks) while adhering to the specified constraints. There is by now a very rich academic literature on the design of autobidding algorithms for various constraints and auction designs [8, 9, 11, 12, 14, 17] and how a platform might design its auction rules or prices for such algorithms [7, 24, 29].

From the advertisers' perspective, the autobidder and auction layers can be seen as a cohesive mechanism that the advertisers interact with strategically. The advertiser-specified constraints (e.g., budgets) define the mechanism's message space. However, this is not a direct-revelation mechanism, and the allowable constraints do not necessarily capture the intricacies of advertiser preferences. Indeed, the latter may be much more complex than, e.g., maximizing a linear value for clicks or other events subject to a budget.[2]

Thus, while the autobidders play a bidding game amongst themselves, the advertisers play a game of constraint selection, henceforth called the *metagame*. The actions in this metagame are the budgets and any other constraints specified to the autobidders. These constraints in turn define the bidding game, called the *inner*

[1]In our model, maximum bid constraints are equivalent to average return-on-investment (ROI) constraints, which are common in autobidding. See Remark 2.5.

[2]And even if the message space does capture advertisers' preferences, these mechanisms do not generally incentivize truth-telling; see, e.g., [3].

STOC '24, June 24–28, 2024, Vancouver, BC, Canada
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0383-6/24/06
https://doi.org/10.1145/3618260.3649688

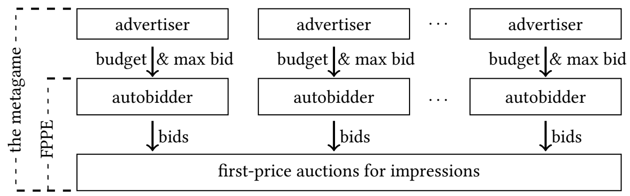

Figure 1: The metagame of strategic budget selection.

game, played by the autobidders. The outcome of the inner game ultimately determines the advertisers' payoffs. In this paper we focus on the metagame played by the advertisers, explore its equilibrium outcomes, and analyze the efficiency of the market.

Our model. For the inner game, each autobidder receives as input a budget constraint and maximum bid, either of which can be infinite. We first need to specify the outcome of the inner game between the autobidders for a fixed choice of constraints. One potential challenge is that if there are multiple equilibria in the inner game, the metagame payoffs would be ambiguous and dependent on the details of the learning dynamics. Fortunately, if the auction format is a first-price auction, the inner game has an essentially unique equilibrium known as the *first-price pacing equilibrium (FPPE)* [20]. Moreover, common learning dynamics are known to converge quickly to this equilibrium [14]. Given that and the practical prevalence of first-price auctions in the online advertising market, we will focus on first-price auctions in our model and use FPPE as our predicted outcome of the inner game.[3] This approach allows us to abstract away from the details of the autobidder implementation; our results are relevant to any learning methods that converge to the FPPE.

Our model encompasses a broad range of advertiser preferences. We consider general *separable* utility models with a weakly concave (not necessarily linear) valuation for events (e.g., clicks) and weakly convex increasing disutility for spending money. This includes quasi-linear utility up to a hard budget constraint, but also allows value-maximizing advertisers, soft budget constraints, decreasing marginal value for clicks, etc. We emphasize the distinction between preferences and constraints: while the autobidders strive to maximize value given the specified constraints, the advertisers are optimizing for their more general preferences.

To measure the efficiency of the overall market, we focus on the *liquid welfare*, which is the summation of each advertiser's willingness-to-pay given her own type and allocation. This general definition encompasses the classic definition of welfare for linear agents and the original definition of liquid welfare for budgeted agents, respectively [25, 43]. Liquid welfare is also closely related to *compensating variation* from economics [31].[4] We emphasize that liquid welfare is determined by the allocation and the advertisers' true preferences, not the payments or the declared constraints. In our results, we compare to the optimal liquid welfare achieved by *any* allocation, henceforth called the OPT.

Approximate efficiency in the metagame. In our baseline model, we consider a full-information environment where each advertiser's type (i.e., valuation and disutility for spending money) is fixed and publicly known. We first consider pure Nash equilibria of the metagame: the advertisers simultaneously declare their constraints to their respective autobidders then outcomes and payoffs are determined by the FPPE of the inner game.

While the FPPE is essentially unique given the choices of the advertisers, the same is not true for the metagame of constraint selection. We show through a sequence of examples that the metagame may have multiple pure Nash equilibria, or none at all. Intuitively, a unilateral change of budget by one advertiser can substantially change the equilibrium behavior of *all* autobidders in the inner game, leading to a rich strategic environment for the advertisers. Nevertheless, in our main result we prove that at any pure Nash equilibrium (if one exists), the resulting liquid welfare is a 2-approximation to the OPT and this is tight.[5]

We also show that at any mixed Nash equilibrium (which is guaranteed to exist), the resulting expected liquid welfare is at most a 4-approximation to the OPT. In fact, this approximation result extends to Bayesian environments, where each advertiser's private type is independently drawn from a publicly-known prior. In this extended setting, we can likewise show that the expected liquid welfare at any Bayesian Nash equilibrium is at most a 4-approximation to the expected OPT.

We reiterate that our approximation bounds hold for any (distribution over) separable preferences of the advertisers.[6] We view the generality of our bounds as a benefit of the design paradigm of mechanisms with proxy autobidders, which provide an interface for advertisers with complex preferences to engage with a simple and robust auction format.

Importance of budget constraints. Our approximation results highlight an important practical implication: *Simple instructions for autobidders are sufficient to achieve good market outcomes even when advertisers have much more complex preferences.* An interesting follow-up question arises: Can these instructions be further simplified? Is it necessary to provide advertisers with the option to specify both hard budgets and maximum bids? To answer this, we study variants of the metagame where advertisers exclusively declare budget constraints or maximum bid constraints, respectively. We demonstrate that the budget constraints are crucial. If advertisers specify maximum bids but not budgets, the resulting approximation factor at equilibrium can be as bad as linear in the number of advertisers. On the other hand, all our approximation results extend to the variant of the metagame with no maximum bid.

Intuitively, the presence of budgets allows advertisers to hedge against unexpectedly high expenditures, which can have a potentially catastrophic impact on their utility.

[3]In contrast, the inner game equilibrium may not be unique for 2nd-price auctions [21], and finding any equilibrium is PPAD-hard [18]. So extending to 2nd-price auctions necessitates taking a stance on equilibrium selection and/or the learning dynamics, which we leave for future work.

[4]Compensating variation is the transfer of money required, after some market change, to return agents to their original utility levels. In our scenario, the "change" is an allocation of resources, and the compensating variation is the maximum amount the agents would pay for that allocation; i.e., the liquid welfare.

[5]We say that liquid welfare at equilibrium is an α-approximation to the OPT, for some approximation factor $\alpha \geq 1$, if the liquid welfare is at least OPT$/\alpha$ for any vector of advertiser types.

[6]To the best of our knowledge, the only other $O(1)$-approximation for such separable utilities in a multi-item auction is due to [6], who study direct bidding in second-price auctions; their bound applies to pure Nash equilibria under an additional ROI-optimality refinement and linear valuations. See Section 1 for further discussion.

Characterization of equilibria for homogeneous items. To shed light on the equilibrium structure, we also explicitly characterize the pure Nash equilibria in the special case where all ad impressions are homogeneous. We show that a pure Nash equilibrium always exists in this homogeneous case and can be found in polynomial time. In fact, there are two distinct *types* of equilibria, which we call low-price equilibria and high-price equilibria. There may be multiple equilibria of each type, corresponding to a range of implementable per-unit prices. In a low-price equilibrium, all advertisers specify budget constraints that target their minimal utility-maximizing allocations at a market-clearing price. In a high-price equilibrium, one or more advertisers forego budget constraints and instead impose high maximum bids, effectively forcing a higher price. This multiplicity demonstrates the rich strategic landscape of the metagame, even in the special case where all ad impressions are identical.

Summary of Techniques. Let us now return to our approximation results and provide some additional intuition into the proof ideas. At any pure Nash equilibrium, each advertiser will face a Pareto curve that represents the tradeoff between the value she receives and the budget constraint she declares to her autobidder. Increasing the declared budget allows the advertiser to win more, but at a potentially decreasing bang-per-buck. The first step of our analysis is to analyze this tradeoff. To that end, we will employ a useful interpretation of FPPE due to [20]: the inner FPPE of autobidders corresponds to a market-clearing outcome that assigns a price to each impression, so that each autobidder obtains a preferred bundle under those prices. Since increasing one autobidder's budget has downstream effects on how other autobidders will behave (even keeping the specified constraints of other advertisers fixed), adjusting one's budget can cause the market-clearing prices to change. This gives us the perspective we need to understand an advertiser's tradeoff when choosing a budget: the sensitivity of FPPE prices and market-clearing allocations to budget changes.

Unfortunately, this sensitivity can be unbounded in general. An advertiser might need to increase their spend by an arbitrary amount to secure a target increase in allocation. For example, consider two advertisers competing for a single impression type. If the advertisers set finite budgets and no maximum bids, then at FPPE they will spend their budgets exactly and split the item in proportion to their declared budgets. (See Section 3 for examples and intuition for FPPE.) This means that if the first advertiser sets a budget of 1, the second advertiser would need to set a budget of 9 to receive 90% of the item (at a price of 10 per unit), but a budget of 99 to receive 99% of the item (at a price of 100 per unit). In such a scenario, large changes in price might be required to implement a small change in allocation.

However, a key insight is that this hyper-inflation only occurs when one advertiser obtains a very large allocation of an impression type. As long as an advertiser is not winning an excessive share of the impressions, the market-clearing prices faced by her autobidder will not be overly sensitive to small changes in her allocation. This is our main technical lemma, Lemma 4.2, which relates the rate of substitution between valuation utility and payment disutility of an advertiser at equilibrium to the FPPE prices and the share of each impression she obtains. It implies that if an advertiser is obtaining significantly less value than she would in the optimal allocation, then either (a) the prices are low and her allocation is small, in which case by Lemma 4.2 she can improve her utility by increasing her budget (contradicting the equilibrium assumption), or (b) her autobidder must be facing high market-clearing prices. But in the latter case, we can employ a standard trick for bounding the price of anarchy: the lost liquid welfare from the advertiser can be charged against revenue collected from the other advertisers who are paying these high prices. Here we use the fact that since liquid welfare measures total willingness to pay, the revenue collected is always a lower bound on liquid welfare. Putting these pieces together yields our 2-approximation result for pure Nash equilibrium.

Organization. We formalize the model and provide necessary preliminaries and notations in Section 2. In Section 3 we study two simple examples and analyze their pure Nash equilibrium in the metagame. In Section 4, we present the approximation results of the liquid welfare at equilibrium. We study the extension of our model in the Bayesian environments in Section 5, and several variants of our model under restrictive message space (i.e., advertisers can only exclusively declare hard budgets, or maximum bids) in Section 6. Finally, in Section 7 we characterize and prove existence of pure Nash equilibria for single-item (aka., homogeneous) instances. All missing proofs are deferred to the full version.

Below we discuss the additional related work.

Incentives in Autobidding. A recent line of work has studied the incentives of value-maximizing advertisers to truthfully describe their preferences to their autobidders. [40] show that the advertisers may benefit by misreporting their preference to autobidders when the platform does not have commitment power and can change the auction rule. [35] conduct an empirical investigation using numerical examples, illustrating the potential advantages for advertisers when misreporting to autobidders in a second-price auction. Conceptually closest to our work is [3], who introduce the concept of auto-bidding incentive compatibility (AIC) and show that first-price auction with uniform bidding satisfies AIC (while many other auction formats do not). Crucially, all of these works assume that both advertisers and their autobidders are value-maximizers with similar constraints, so the advertisers are able to declare their true preferences to their autobidders. In contrast, in our model the advertisers are utility-maximizers and they can only give their autobidders simple constraints (hard budget, maximum bid) that may not be rich enough to capture their true complex preferences. In particular, the induced mechanism is not direct-revelation and hence incentive compatibility does not apply; we focus instead on Nash equilibria.

Metagame between no-regret learning agents. [34] explores the concept of a metagame involving rational agents who utilize regret-minimizing learning algorithms to play games on their behalf, and investigates whether the agents are incentivized to manipulate or misrepresent their true preferences in various classic games. In [33], the authors delve into a similar problem but within the context of auctions involving two agents with linear utility, showing that agents have incentive to misreport their true values in a second-price auction when both agents use multiplicative-weights learning algorithms. On the other hand, they find that truthful reporting forms a Nash equilibrium in the first-price auction when both agents use mean-based learning algorithms. Compared to both [34] and

[33], this present paper considers a different auction setup with budget constraints, focuses on cases where advertiser preferences and autobidder instructions are different, and abstracts away from the learning process by assuming that the autobidders converge to bidding according to an FPPE of the simultaneous auction game.

Price of anarchy for non-quasi linear agents. There is a long line of literature about the price of anarchy (PoA) – the approximation between worst equilibrium and best outcome – in the context of liquid welfare for non-quasi linear agents. Besides works already mentioned above, [25] introduce the concept of liquid welfare for agents with hard budget constraints. They prove the PoA for posting market clearing prices and the clinching auction. [37] design a sampling mechanism with a better PoA guarantee. All mechanisms studied in [25, 37] are truthful mechanisms. For non-truthful mechanisms, [15, 16, 19] study the PoA of the simultaneous Kelly mechanism, while [5] studies the PoA of the simultaneous first-price auction and second-price auction with no over-bidding. [6] studies the PoA under ROI-optimal pure Nash equilibrium of the simultaneous second price auction, establishing a 2-approximation and that such a pure Nash equilibrium always exists. [1, 22, 23, 36, 39] study the PoA for the autobidders under various auctions. [27, 28, 38] study the dynamic of no-regret learning/budget-pacing players (autobidders) and provide the liquid welfare guarantees. In this literature [1, 23, 27, 28, 36, 38], similar 2-approximation PoA results are obtained for autobidders in various auction formats for both static and dynamic environments. It is important to highlight that these works treat the autobidder constraints as exogenous and do not model them as strategic choices, focusing rather on the auction and autobidders' interaction. Additionally, except [6], the aforementioned studies assume that agents have a linear disutility for spending money up to a fixed budget. In contrast, our paper (like 6) considers agents with a general convex disutility function for spending money with a hard budget.

2 MODEL AND PRELIMINARIES

Agent models. There are $n \geq 2$ agents (advertisers) and m divisible items (impressions). The outcome for agent i is (x_i, t_i), where $x_i = (x_{i1}, \dots, x_{im}) \in [0,1]^m$ is the allocation for each item j and $t_i \in \mathbb{R}_+$ is the payment. Given allocation x_i, agents receives $\sum_{j\in[m]} \phi_{ij}x_{ij}$ number of clicks, where $\phi_{ij} \in \mathbb{R}_+$ is the *click-through rate* of each item j for agent i. The click-through rate can also be interpreted as conversion rate or other related concepts in different applications.

Agent i's von Neumann–Morgenstern utility is parameterized by her *type* (V_i, w_i, C_i): Given outcome (x_i, t_i), agent i's utility $u_i(x_i, t_i)$ is defined as

$$u_i(x_i, t_i) \triangleq \begin{cases} V_i\left(\sum_{j\in[m]} \phi_{ij}x_{ij}\right) - C_i(t_i) & \text{if } t_i \leq w_i \\ -\infty & \text{if } t_i > w_i \end{cases}$$

where $V_i : \mathbb{R}_+ \to \mathbb{R}_+$ is the *valuation function* mapping from the total number of received clicks to agent's valuation; $w_i \in \mathbb{R}_+^\infty$ is the *hard budget*;[7] and $C_i : \mathbb{R}_+ \to \mathbb{R}_+$ is the *money cost function* mapping from the payment to the disutility for spending money.[8] We assume V_i is differentiable, weakly concave, weakly increasing, and $V_i(0) = 0$; and money cost function C_i is differentiable, weakly convex, weakly increasing, and $C_i(0) = 0$. We will also assume that either $w_i < \infty$ or $C_i(t_i) > 0$ for some t_i. This rules out agents with no value for money or spending constraint. We write $S_i(\sum_j \phi_{ij}x_{ij})$ and $R_i(t_i)$ as the derivative of V_i and C_i at $\sum_j \phi_{ij}x_{ij}$ and t_i, respectively. Agents with this utility model are called *general agents*.

Three classic models can be viewed as special cases:

(1) *(linear utility)* An agent i with linear utility with type $v_i \in \mathbb{R}_+$, hereafter *linear agent*, has linear valuation function $V_i(\sum_j \phi_{ij}x_{ij}) = v_i \cdot \sum_j \phi_{ij}x_{ij}$ where v_i is her *value per click*, no hard budget (aka., $w_i = \infty$), and identity money cost function $C_i(t_i) = t_i$.

(2) *(budgeted utility)* An agent i with budgeted utility with type $(v_i, w_i) \in \mathbb{R}_+ \times \mathbb{R}_+^\infty$, hereafter *budgeted agent*, has linear valuation function $V_i(\sum_j \phi_{ij}x_{ij}) = v_i \cdot \sum_j \phi_{ij}x_{ij}$ where v_i is her *value per click*, hard budget w_i, and identity money cost function $C_i(t_i) = t_i$.

(3) *(value-maximizing utility)* A *value-maximizing* agent i with type $(v_i, w_i) \in \mathbb{R}_+^2$, has linear valuation $V_i(\sum_j \phi_{ij}x_{ij}) = v_i \cdot \sum_j \phi_{ij}x_{ij}$ where v_i is her *value per click*, hard budget w_i, and zero money cost function $C_i(t_i) = 0$ for $t_i \leq w_i$. Value-maximizing agents have been studied extensively in the recent autobidding literature [10]. Clearly, every linear agent is a budgeted agent, and every budgeted agent is a general agent.

First-price pacing equilibrium for budgeted agents (autobidders). Before introducing our solution concept, we first revisit a pivotal concept: the *first-price pacing equilibrium*, introduced by [20] to study budgeted agents (autobidders) in an advertising market.[9]

Definition 2.1 (First-price pacing equilibrium). *For budgeted agents with types $\{(v_i, w_i)\}_{i\in[n]}$, a* first-price pacing equilibrium (FPPE) *is a tuple (p, x, t) of per-unit price $p_j \in \mathbb{R}_+$ for each item j, allocation $x_i \in [0,1]^m$, and payment $t_i \in \mathbb{R}_+$ for each agent i that satisfies the following properties: (highest bang-per-buck) if $x_{ij} > 0$, then $j \in \operatorname{argmax}_{j'\in[m]} \frac{v_i\phi_{ij'}}{p_{j'}}$ and $v_i\phi_{ij} \geq p_j$; (supply feasibility) $\sum_{i\in[n]} x_{ij} \leq 1$ and equality holds if $p_j > 0$; (payment calculation) $t_i = \sum_{j\in[m]} x_{ij}p_j$; (budget feasibility) $t_i \leq w_i$ and equality holds if $\max_{j\in[m]} \frac{v_i\phi_{ij}}{p_j} > 1$.*

As noted by [20], the per-unit prices, allocation, and payments in FPPE can be interpreted as the outcome of a budget-pacing game for first-price auctions, as follows. Each budgeted agent (autobidder) i first determines an item-independent *pacing multiplier* $\alpha_i = 1 \wedge (\min_{j\in[m]} \frac{p_j}{v_i\phi_{ij}}) \in [0,1]$. The agent then submits bid $\alpha_i v_i \phi_{ij}$ for each item j.[10] Then the allocation is the highest-bids-win and the payment is computed under the first-price format. As a sanity check,

[7] We use notation $\mathbb{R}_+^\infty$ to denote the set of all non-negative real numbers and infinite, i.e., $\mathbb{R}_+^\infty = \mathbb{R}_+ \cup \{\infty\}$.

[8] Money cost function C_i is interpreted as follows: Agent i has an outside option value for her money, e.g., spending on another advertising platform. Then $C_i(t_i)$ is the utility that buyer i foregoes by paying t_i to the current platform (seller).

[9] For ease of presentation, we use an equivalent definition of the first-price pacing equilibrium. See the full version for the original definition from [20] and the proof of equivalence.

[10] Here the goal of each autobidder is to maximize the total value (or equivalently total number of clicks due to the linear valuation function) received subject to the budget constraint. Under this interpretation, the pacing multiplier can be considered as the Lagrangian multiplier of the budget constraint.

under this interpretation, the "highest bang-per-buck" property is satisfied due to the definition of pacing scalar α_i and the highest-bids-win allocation construction.

The FPPE can also be interpreted as a (supply-unaware) competitive equilibrium: if we fix the per-unit price p_j for each item j, the allocation x_i maximizes the utility of agent i subject to her budget w_i, without accounting for the supply feasibility. We will occasionally refer to this interpretation.

Lemma 2.1 (Conitzer et al. [20]). *For any set of budgeted agents, FPPE exists. Moreover, the per-unit price of each item and utility of each budgeted agent are unique.*

FPPE is the solution concept of interest for the autobidders, due to its existence and uniqueness guarantee. Moreover, common learning dynamics (e.g., dynamic competition of autobidders) in repeated first-price auctions are known to converge quickly to this equilibrium when budgeted agents use linear strategies (i.e., identical pacing multipliers for all items) [13]. We denote $p(\{(v_i, w_i)\}_{i\in[n]})$, $x^\sigma(\{(v_i, w_i)\}_{i\in[n]})$, and $t^\sigma(\{(v_i, w_i)\}_{i\in[n]})$ by the unique per-unit prices and the corresponding allocation, payment (under tie-breaking rule σ) in the FPPE for budgeted agents with types $\{(v_i, w_i)\}_{i\in[n]}$, respectively.[11]

Let's present two examples to illustrate the concept of FPPE; we reuse them in Section 3.

Example 2.2 (Two linear agents and single item). *Consider two linear agents and one item, such that the click-through rates are the same for both agents, i.e., $\phi_{11} = \phi_{21} = 1$. Further, the value per click v_1 for agent 1 is greater than or equal to the value per click v_2 for agent 2, i.e., $v_1 \ge v_2$.*

In the FPPE, the unique per-unit price for the item is $p_1 = v_1$. If $v_1 > v_2$, the allocation in the FPPE is also unique, with $x_{11} = 1$ and $x_{21} = 0$. On the other hand, if $v_1 = v_2$, any allocation $x_{11}, x_{21} \in [0, 1]$ satisfying $x_{11} + x_{21} = 1$, along with the unique per-unit price $p_1 = v_1$, forms an FPPE.

Example 2.3 (Two budgeted agents and two items). *Suppose there are two budgeted agents and two items. Let us assume that the click-through rate ϕ_{ij} is given by $\frac{1}{2} + \frac{1}{2}\mathbb{I}\{i = j\}$ for each $i \in [2]$ and $j \in [2]$, i.e., each agent i favors item i than the other item. The value per click is the same for both agents, i.e., $v_1 = v_2 = 1$. Both agents have a budget constraint of $w_1 = w_2 = \frac{1}{2}$.*

In the FPPE, both the per-unit price and the allocation are unique. Specifically, we have $p_1 = p_2 = \frac{1}{2}$, and the allocation $x_{ij} = \mathbb{I}\{i = j\}$ for each $i \in [2]$ and $j \in [2]$. In other words, both agents receive their favorite items as per the allocation.

Metagame of strategic budget selection for general agents. An FPPE is defined only for budgeted agents. For more general agents, we will take inspiration from autobidding platforms in practice and imagine the agents are provided an interface to report a budgeted agent's type that will specify the behavior of an autobidder. This defines a *metagame of strategic budget selection* for agents with general utility models:

Definition 2.4 (Metagame of strategic budget selection). *Each general agent i decides on a message $(\tilde{v}_i, \tilde{w}_i) \in (\mathbb{R}_+^\infty)^2$ as her report to the seller.*[12] *Given reported message profile $\{(\tilde{v}_i, \tilde{w}_i)\}_{i\in[n]}$, the seller implements allocation $x(\{(\tilde{v}_i, \tilde{w}_i)\}_{i\in[n]})$ and payment $t(\{(\tilde{v}_i, \tilde{w}_i)\}_{i\in[n]})$ induced by the FPPE, assuming that agents have budgeted utility with types $\{(\tilde{v}_i, \tilde{w}_i)\}_{i\in[n]}$.*

Going forward, we will refer to the metagame of budgeted utility reporting as simply the *metagame* for the sake of brevity. Similarly, the FPPE induced by a message profile will be referred to as the *inner FPPE.* Given the structure of FPPE, the message $(\tilde{v}_i, \tilde{w}_i)$ reported by agent i in the metagame can also be interpreted as the constraints on the maximum bid $\tilde{v}_i$ and constraints on the maximum payment (hard budget) $\tilde{w}_i$ that agent i specifies to her autobidder. The agent has the option to report $\tilde{v}_i = \infty$ ($\tilde{w}_i = \infty$), indicating the absence of any constraint on the maximum bid or maximum payment, respectively.

Remark 2.5 (Return on Investment Constraints). *Another common type of autobidding constraint is an (aggregate) ROI constraint, which bounds the ratio between the total number of allocated clicks and the total payment. We note that, for FPPE, a maximum bid $\tilde{v}_i$ is equivalent to an aggregate ROI constraint. Indeed, due to the "highest-bang-per-buck" property of FPPE, all items allocated to an autobidder have the same ROI, which is equal to the equilibrium bid. So an ROI constraint of the form $t_i \le \gamma \sum_j \phi_{ij} x_{ij}$ excludes equilibrium bids higher than γ, and a maximum bid of $\tilde{v}_i$ guarantees an average payment of at most $\tilde{v}_i$ per click.*

With slight abuse of notations, we use $u_i(\tilde{v}_i, \tilde{w}_i, \tilde{v}_{-i}, \tilde{w}_{-i})$, $x_i(\tilde{v}_i, \tilde{w}_i, \tilde{v}_{-i}, \tilde{w}_{-i})$, $t_i(\tilde{v}_i, \tilde{w}_i, \tilde{v}_{-i}, \tilde{w}_{-i})$ to represent the utility, allocation and payment of agent i in the metagame when agent i reports message $(\tilde{v}_i, \tilde{w}_i)$ and other agents report message $(\tilde{v}_{-i}, \tilde{w}_{-i}) \triangleq \{(\tilde{v}_{i'}, \tilde{w}_{i'})\}_{i' \neq i}$.[13] Similarly, we use $p_j(\tilde{v}_i, \tilde{w}_i, \tilde{v}_{-i}, \tilde{w}_{-i})$ to represent the per-unit price of item j of the inner FPPE given message profile $\{(\tilde{v}_i, \tilde{w}_i), (\tilde{v}_{-i}, \tilde{w}_{-i})\}$. The following lemma characterizes the relationship between agents' utility, payment, and per-unit prices. Essentially, for each agent, her payment along with the per-unit prices of the inner FPPE serve as sufficient statistics for computing her utility. Its proof is straightforward given the definitions of FPPE and metagame.

Lemma 2.2. *In the metagame, for every agent i with type $\{V_i, w_i, C_i\}$ and every message profile, let t_i be the payment of agent i and p be the per-unit prices of the inner FPPE. Then agent i's utility u_i satisfies*

$$u_i = \begin{cases} V_i\left(\left(\max_{j\in[m]} \frac{\phi_{ij}}{p_j}\right) \cdot t_i\right) - C_i(t_i) & \text{if } t_i \le w_i \\ -\infty & \text{if } t_i > w_i \end{cases}$$

Equilibria in the metagame. In the base model, we are interested in Nash equilibria for agents with public types.

Definition 2.6. *Consider agents with types $\{(V_i, w_i, C_i)\}_{i\in[n]}$. A* pure Nash equilibrium *is a message profile $\{(\tilde{v}_i, \tilde{w}_i)\}_{i\in[n]}$ such that for every agent i and every message $(\tilde{v}_i^\dagger, \tilde{w}_i^\dagger)$,*

$$u_i(\tilde{v}_i, \tilde{w}_i, \tilde{v}_{-i}, \tilde{w}_{-i}) \ge u_i(\tilde{v}_i^\dagger, \tilde{w}_i^\dagger, \tilde{v}_{-i}, \tilde{w}_{-i}).$$

[11]We assume that the seller decides a tie-breaking rule σ exogenously ex ante. In real-world applications, the autobidders are expected to effectively "implement" such rule by making micro-adjustments to their bids across time, in order to hit their budgets in aggregate. Our single-shot model abstracts away from this dynamic behavior, but it motivates our focus on FPPE allocations. Importantly, our results are independent of the specific choice of σ. We omit mentioning σ when it is clear from the context.

[12]We use "$\tilde{\ }$" to denote the budgeted utility model that an agent reports to the seller.

[13]We use notation $-i$ to denote other $n - 1$ agents excluding agent i.

A mixed Nash equilibrium *is a randomized message profile* $\{(\tilde{\mathbf{v}}_i, \tilde{\mathbf{w}}_i)\}_{i\in[n]}$ *such that*[14] *the random messages are mutually independent, and for every agent i and every message* $(\tilde{v}_i^\dagger, \tilde{w}_i^\dagger)$,

$$\mathbb{E}_{(\tilde{v}_i,\tilde{w}_i)\sim(\tilde{\mathbf{v}}_i,\tilde{\mathbf{w}}_i),(\tilde{v}_{-i},\tilde{w}_{-i})\sim(\tilde{\mathbf{v}}_{-i},\tilde{\mathbf{w}}_{-i})}\left[u_i(\tilde{v}_i,\tilde{w}_i,\tilde{v}_{-i},\tilde{w}_{-i})\right] \geq \mathbb{E}_{(\tilde{v}_{-i},\tilde{w}_{-i})\sim(\tilde{\mathbf{v}}_{-i},\tilde{\mathbf{w}}_{-i})}\left[u_i(\tilde{v}_i^\dagger,\tilde{w}_i^\dagger,\tilde{v}_{-i},\tilde{w}_{-i})\right]$$

As a sanity check, in both pure/mixed Nash equilibrium, the utility of every agent i is non-negative, since zero utility is always guaranteed by reporting message $(\tilde{v}_i = 0, \tilde{w}_i = 0)$.

Liquid welfare. We evaluate a particular allocation in terms of its *liquid welfare.*

Definition 2.7 (Liquid welfare)**.** *For agents with types* $\{(V_i, w_i, C_i)\}_{i\in[n]}$*, the* liquid welfare *of a (possibly) randomized allocation* $\mathbf{x}$ *is*

$$W(\mathbf{x}) \triangleq \sum_{i\in[n]} W_i(\mathbf{x}_i),$$

where $W_i(\mathbf{x}_i)$ *is agent i's* willingness to pay *for allocation* $\mathbf{x}_i$,

$$W_i(\mathbf{x}_i) \triangleq \min\left\{w_i,\ C_i^{-1}\left(\mathbb{E}_{x_i\sim\mathbf{x}_i}\left[V_i\left(\sum_{j\in[m]}\phi_{ij}x_{ij}\right)\right]\right)\right\}.$$

Within the set of feasible randomized allocations, the optimal allocation that maximizes the liquid welfare is deterministic, thanks to the weak concavity of the valuation function V_i.

We evaluate the metagame via an approximation of liquid welfare. Specifically, we compare liquid welfare of the worst equilibria against that of the best allocation, over all instances.[15]

Definition 2.8 (Price of anarchy)**.** *The* price of anarchy (PoA) *of the metagame* $\Gamma_{\texttt{Pure}}$ *(resp.,* $\Gamma_{\texttt{Mixed}}$*) under pure (resp., mixed) Nash equilibrium is*

$$\Gamma_{\texttt{Pure}} \triangleq \sup_{n,m,\phi}\ \sup_{\{V_i,w_i,C_i\}_{i\in[n]}} \frac{\max_x W(x)}{\inf_{x\in\texttt{Pure}} W(x)},$$

$$\Gamma_{\texttt{Mixed}} \triangleq \sup_{n,m,\phi}\ \sup_{\{V_i,w_i,C_i\}_{i\in[n]}} \frac{\max_x W(x)}{\inf_{\mathbf{x}\in\texttt{Mixed}} W(\mathbf{x})}$$

where Pure *(resp.,* Mixed*) is the set of the deterministic (resp., randomized) allocation profile induced by all pure (mixed) Nash equilibrium given types* $\{V_i, w_i, C_i\}_{i\in[n]}$.

3 EXAMPLES: PURE NASH EQUILIBRIUM

Let us revisit Examples 2.2 and 2.3 and analyze pure Nash equilibria therein. We establish:

Proposition 3.1. *In the metagame, pure Nash equilibria might not exist, even for budgeted agents. When they do exist, the resulting per-unit prices may not be unique, even for linear agents, due to the multiplicity of equilibria.*

We interpret this proposition as follows. First, although the inner FPPE of the metagame always exists and its induced per-unit prices are unique (as per Lemma 2.1), incorporating the strategic behavior of the advertisers in the meta-game complicates the allocation and payment outcome. Second, due to the non-existence result, even budgeted agents may have incentive to misreport their types.[16] Finally, despite the non-existence result, we show in Section 7 that a pure Nash equilibrium always exists for single-item instances and can be computed in polynomial time.

We start with two auxiliary lemmas on verifying the existence of profitable deviations in a pure Nash equilibrium. The first lemma suggests that for a given agent i, it suffices to consider deviations $(\tilde{v}_i^\dagger, \tilde{w}_i^\dagger)$ with a restriction that $\tilde{v}_i^\dagger = \infty$, i.e., no maximum bid constraint for her autobidder. Loosely speaking, restricting to deviation with $\tilde{v}_i^\dagger = \infty$ simplifies the analysis since the agent would exhaust her reported budget $\tilde{w}_i^\dagger$. So, Lemma 2.2 ensures that $\tilde{w}_i^\dagger$ along with the per-unit prices of the inner FPPE serve as sufficient statistics for computing her utility.

Lemma 3.2. *In the metagame, for every agent i and every pure Nash equilibrium* $\{(\tilde{v}_i, \tilde{w}_i), (\tilde{v}_{-i}, \tilde{w}_{-i})\}$,

$$u_i(\tilde{v}_i,\tilde{w}_i,\tilde{v}_{-i},\tilde{w}_{-i}) = \max_{\tilde{v}_i^\dagger,\tilde{w}_i^\dagger} u_i(\tilde{v}_i^\dagger,\tilde{w}_i^\dagger,\tilde{v}_{-i},\tilde{w}_{-i}) = \max_{\tilde{w}_i^\dagger} u_i(\infty,\tilde{w}_i^\dagger,\tilde{v}_{-i},\tilde{w}_{-i}).$$

The second lemma suggests that, for a given agent i, it suffices to consider the tie-breaking rule of the inner FPPE that favors agent i. Its proof is straightforward given Lemma 3.2.

Lemma 3.3. *In the metagame with tie-breaking rule* σ *for the inner FPPE, for every agent i and every pure Nash equilibrium* $\{(\tilde{v}_i, \tilde{w}_i), (\tilde{v}_{-i}, \tilde{w}_{-i})\}$*, it satisfies that*

$$u_i^\sigma(\tilde{v}_i,\tilde{w}_i,\tilde{v}_{-i},\tilde{w}_{-i}) = \max_{\sigma'} u_i^{\sigma'}(\tilde{v}_i,\tilde{w}_i,\tilde{v}_{-i},\tilde{w}_{-i})$$

where $u_i^\sigma(\tilde{v}_i,\tilde{w}_i,\tilde{v}_{-i},\tilde{w}_{-i})$ *and* $u_i^{\sigma'}(\tilde{v}_i,\tilde{w}_i,\tilde{v}_{-i},\tilde{w}_{-i})$ *are agent i's utility in the metagame when tie-breaking rules* $\sigma, \sigma^\dagger$ *are selected for the inner FPPE, respectively; and the maximization on the right-hand side is taken over all randomized and deterministic tie-breaking rules.*

3.1 Example 2.2: Non-Uniqueness

We revisit the linear agents instance from Example 2.2. We prove that there exists an efficient pure Nash equilibrium (Claim 3.4), and also an inefficient pure Nash equilibrium if v_2 is close to v_1 (Claim 3.5). In both equilibria, agents report their value per click (aka., maximum bid) truthfully, while strategically declaring their budgets.

Claim 3.4. *In Example 2.2, a pure Nash equilibrium of the metagame is achieved when agent 1 reports* $\tilde{v}_1 = v_1$ *and* $\tilde{w}_1 = v_2$*, while agent 2 reports* $\tilde{v}_2 = v_2$ *and* $\tilde{w}_2 = v_2$.

Proof. In the following argument, we assume $v_1 > v_2$. However, the same argument can be applied in the case of $v_1 = v_2$ due to Lemma 3.3. To save space, we will omit the details of this case.

Given message profile $((\tilde{v}_1 = v_1, \tilde{w}_1 = v_2), (\tilde{v}_2 = v_2, \tilde{w}_2 = v_2))$, the inner FPPE allocates the entire item to agent 1 with per-unit price v_2, i.e., $x_{11} = 1$, $x_{21} = 0$ and $p_1 = v_2$.

We verify the non-existence of profitable deviation. Invoking Lemma 3.2, it is sufficient to consider deviation $(\tilde{v}_i^\dagger, \tilde{w}_i^\dagger)$ with $\tilde{v}_i^\dagger =$

[14]We use bold symbols (e.g., $\tilde{\mathbf{v}}, \tilde{\mathbf{w}}$) to denote random variables and their corresponding distributions.

[15]Though the inner FPPE of the metagame assumes that each autobidder uses a single pacing multiplier and conducts "linear bidding", the price of anarchy compares its equilibrium efficiency with the unrestricted optimal allocation. Our approximation results hold in spite of the restriction to linear bidding in the inner FPPE.

[16]For this point, it is crucial that the budgeted agents are utility-maximizers. Indeed, in the meta-game for budgeted *value*-maximizing agents, truthful reporting is a Nash equilibrium (this follows from 3).

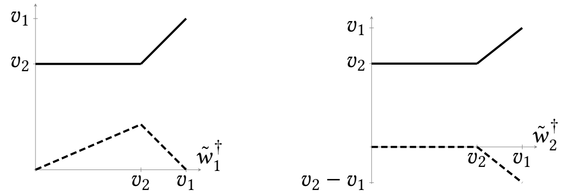

(a) deviation of agent 1 (b) deviation of agent 2

Figure 2: Graphical illustration of each agent i's deviation $(\infty, \tilde{w}_i^\dagger)$ in Claim 3.4 for Example 2.2. The solid (resp., dashed) line is the per-unit price (resp., utility of agent i).

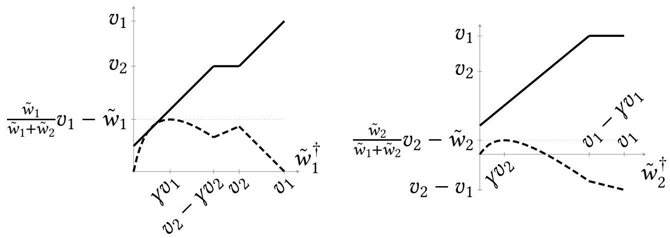

(a) deviation of agent 1 (b) deviation of agent 2

Figure 3: Graphical illustration of each agent i's deviation $(\infty, \tilde{w}_i^\dagger)$ in Claim 3.5 for Example 2.2. The solid (dashed) line is the per-unit price (utility of agent i).

∞ for each agent i. In such a deviation, the "budgeted feasibility" property of FPPE ensures that agent i exhausts her budget and pays exactly $\tilde{w}_i^\dagger$. We now proceed to analyze each agent i individually. For agent 1, it can be verified that the per-unit price $p_1^\dagger$ and allocation $x_{11}^\dagger$ of inner FPPE under her deviation $(\infty, \tilde{w}_1^\dagger)$ are

$$p_1^\dagger = \begin{cases} v_2 & \text{if } \tilde{w}_1^\dagger \le v_2 \\ \tilde{w}_1^\dagger & \text{if } \tilde{w}_2^\dagger \ge v_2 \end{cases}, x_{11}^\dagger = \frac{\tilde{w}_1^\dagger}{p_1^\dagger} = \begin{cases} \frac{\tilde{w}_1^\dagger}{v_2} & \text{if } \tilde{w}_1^\dagger \le v_2 \\ 1 & \text{if } \tilde{w}_2^\dagger \ge v_2 \end{cases}$$

and her utility is maximized at $\tilde{w}_1^\dagger = v_2$, which results in the same utility as she obtains in equilibrium. See Figure 2a for a graphical illustration. To avoid repetition, we omit a similar argument for agent 2. See Figure 2b for a graphical illustration. □

Claim 3.5. *Consider the metagame in Example 2.2 with $\frac{\sqrt{5}-1}{2}v_1 \le v_2 \le v_1$. Let $\gamma = \frac{v_1 v_2}{(v_1+v_2)^2}$. Then a pure Nash equilibrium is achieved when agent 1 reports $\tilde{v}_1 = v_1$ and $\tilde{w}_1 = \gamma v_1$, while agent 2 reports $\tilde{v}_2 = v_2$ and $\tilde{w}_2 = \gamma v_2$.*

Proof. In the following argument, we assume $v_1 > v_2$. However, the same argument can be applied in the case of $v_1 = v_2$ due to Lemma 3.3. To save space, we will omit the details of this case.

Given message profile $((\tilde{v}_1 = v_1, \tilde{w}_1 = \gamma v_1), (\tilde{v}_2 = v_2, \tilde{w}_2 = \gamma v_2))$, the per-unit price p_1 of the inner FPPE is $p_1 = \tilde{w}_1 + \tilde{w}_2 = \gamma(v_1+v_2) \le v_2 < v_1$. Moreover, the item is allocated to both agents in proportion to their respective values, i.e., $x_{i1} = \frac{v_i}{(v_1+v_2)}$.

We verify the non-existence of profitable deviation. Invoking Lemma 3.2, it is sufficient to consider deviation $(\tilde{v}_i^\dagger, \tilde{w}_i^\dagger)$ with $\tilde{v}_i^\dagger = \infty$ for each agent i. In such a deviation, the "budgeted feasibility" property of FPPE ensures that agent i exhausts her budget and pays exactly $\tilde{w}_i^\dagger$. We now proceed to analyze each agent i individually. For agent 1, it can be verified that the per-unit price $p_1^\dagger$ and allocation $x_{11}^\dagger$ of inner FPPE under her deviation $(\infty, \tilde{w}_1^\dagger)$ are

$$p_1^\dagger = \begin{cases} \tilde{w}_1^\dagger + \tilde{w}_2 & \text{if } \tilde{w}_1^\dagger \le v_2 - \tilde{w}_2 \\ v_2 & \text{if } v_2 - \tilde{w}_2 \le \tilde{w}_2^\dagger \le v_2 \\ \tilde{w}_1^\dagger & \text{if } \tilde{w}_2^\dagger \ge v_2 \end{cases},$$

$$x_{11}^\dagger = \frac{\tilde{w}_1^\dagger}{p_1^\dagger} = \begin{cases} \frac{\tilde{w}_1^\dagger}{\tilde{w}_1^\dagger + \tilde{w}_2} & \text{if } \tilde{w}_1^\dagger \le v_2 - \tilde{w}_2 \\ \frac{\tilde{w}_1^\dagger}{v_2} & \text{if } v_2 - \tilde{w}_2 \le \tilde{w}_2^\dagger \le v_2 \\ 1 & \text{if } \tilde{w}_2^\dagger \ge v_2 \end{cases}$$

By considering the first-order condition, we observe that agent 1's utility, under the mentioned deviation, has two local maximizers: $\tilde{w}_1^\dagger = \gamma v_1$ and $\tilde{w}_1^\dagger = v_2$. However, according to the claim assumption that $\frac{(\sqrt{5}-1)v_1}{2} \le v_2$, her utility is maximized at $\tilde{w}_1^\dagger = \gamma v_1$, which coincides with her utility in the equilibrium. See Figure 3a for a graphical illustration. To avoid repetition, we omit a similar argument for agent 2. See Figure 3b for a graphical illustration. □

Remark 3.1. *At the equilibrium of the metagame, each agent i faces a Pareto curve (see Figures 2 and 3) that describes how much value they receive as their reported budget $\tilde{w}_i$ increases. Increasing her reported budget $\tilde{w}_i$ causes more impressions and thus clicks to be won, but at a potentially decreasing bang-per-buck. In particular, as we mentioned in Section 2, FPPE can be viewed as a competitive equilibrium, i.e., market-clearing outcome that assigns a price to each impression. Since increasing budget has downstream effects on how other autobidders will behave (even keeping the reports of other advertisers fixed), increasing one's budget can cause prices to increase.*

3.2 Example 2.3: Non-Existence

We revisit Example 2.3 and prove non-existence of pure Nash equilibrium for budgeted agents.

Claim 3.6. *In the metagame from Example 2.3, pure Nash equilibrium does not exist. Consequently, both agents have incentive to misreport their true types.*

To prove this claim, we enumerate all feasible allocations and argue that each of them cannot be induced by a pure Nash equilibrium. We distinguish three cases with different deviation strategy for each, see Claims 3.7, 3.8 and 3.9. In the first case, we argue that allocations where each agent receives a positive fraction of her less favored item cannot be induced by a pure Nash equilibrium.

Claim 3.7. *In Example 2.3, there exists no pure Nash equilibrium whose induced allocation x satisfies $x_{12} > 0$ and $x_{21} > 0$.*

Proof. We prove this by contradiction. Suppose there exists a pure Nash equilibrium as desired. Let p_1, p_2 be the per-unit price of the inner FPPE. Since $x_{12} > 0$ in FPPE, the "highest bang-per-buck" property implies $\frac{\phi_{12}}{p_2} \ge \frac{\phi_{11}}{p_1}$ and thus $p_2 \le \frac{1}{2}p_1$. Similarly,

$x_{21} > 0$ in FPPE implies $p_1 \le \frac{1}{2}p_2$. Thus, the inner FPPE has zero per-unit prices, i.e., $p_1 = p_2 = 0$, which can only be achieved from message profile $\{(\tilde{v}_i, \tilde{w}_i)\}$ where $\tilde{v}_i = 0$ or $\tilde{w}_i = 0$ for each agent i. It is straightforward to verify that $(\tilde{v}_i^\dagger = \epsilon, \tilde{w}_i^\dagger = \epsilon)$ is a profitable deviation for each agent i with sufficiently small ϵ, which leads to a contradiction. □

In the second case, we argue that the allocation where each agent receives her favored item cannot be induced by a pure Nash equilibrium. Note that this is the allocation induced by FPPE if both agents report their types truthfully.

Claim 3.8. *In Example 2.3, there exists no pure Nash equilibrium whose induced allocation x satisfies $x_{11} = x_{22} = 1$.*

Proof. We prove this by contradiction. Suppose there exists a pure Nash equilibrium as desired. Let p_1, p_2 be the per-unit price of the inner FPPE. Without loss of generality,[17] we assume $p_1 \ge p_2$ and $p_1 > 0$. The utility of agent 1 is $u_1 = 1 - p_1$. Consider the following profitable deviation $(\tilde{v}_i^\dagger = \infty, \tilde{w}_i^\dagger = \frac{1}{2}p_2)$. It can be verified that the per-unit price and allocation of inner FPPE under such a deviation are $p_1^\dagger = \frac{1}{2}p_2$ and $p_2^\dagger = p_2$, and $x_{11} = x_{22} = 1$. Consequently, agent 1's utility under such a deviation is $u_1^\dagger = 1 - \frac{1}{2}p_2 \ge u_1$, which leads to a contradiction. □

In the final case, we argue that allocations where one agent i receives her favored item and a positive fraction of less favored item cannot be induced by a pure Nash equilibrium. At a high-level, we utilize the relation $p_i = 2p_{1-i}$ on per-unit prices of inner FPPE due to the "highest bang-per-buck" property, and $x_{ii} = 1, x_{i,1-i} > 0$. We then argue that depending on the magnitude of p_i, either agent has a profitable deviation.

Claim 3.9. *In Example 2.3, there exists no pure Nash equilibrium whose induced allocation x satisfies $x_{ii} = 1$ and $x_{i,1-i} > 0$ for some agent i.*

4 MAIN RESULTS

In this section, we analyze both pure Nash equilibria and mixed Nash equilibria of the metagame.

4.1 Pure Nash Equilibrium

We first present a tight bound on the price of anarchy under pure Nash equilibrium.

Theorem 4.1. *In the metagame, the price of anarchy under pure Nash equilibrium is $\Gamma_{\texttt{Pure}} = 2$.*

Example 4.1 (Lower bound of PoA under pure Nash equilibrium)**.** *Consider a scenario with two agents (one linear agent and one budgeted agent) and one item. Let us assume that the click-through rates are the same for both agents, i.e., $\phi_{11} = \phi_{21} = 1$. Agent 1 has a budget utility model with type $v_1 = K$ and $w_1 = 1$; while agent 2 has a linear utility model with type $v_2 = 1$. Here we assume K is a sufficiently large constant.*

The optimal liquid welfare is $2 - \frac{1}{K}$. This is achieved through an allocation where budgeted agent 1 receives a $\frac{1}{K}$-fraction of the item, and linear agent 2 receives a $\frac{K-1}{K}$-fraction of the item. By employing a similar argument to the one presented in Claim 3.4, we can be verified that a pure Nash equilibrium is achieved when both agents report their types truthfully: $\tilde{v}_1 = v_1 = K$, $\tilde{w}_1 = w_1 = 1$, $\tilde{v}_2 = v_2 = 1$, and $\tilde{w}_2 = \infty$. In this equilibrium, the per-unit price of the inner FPPE is $p_1 = 1$, and agent 1 receives the entire item. Consequently, the achieved liquid welfare is 1. Letting K approach infinity, the lower bound of PoA under pure Nash equilibrium is obtained as desired.

In the rest of this subsection we prove the upper bound in Theorem 4.1: $\Gamma_{\texttt{Pure}} \le 2$. First, we introduce a characterization of the allocation for each agent and per-unit prices of the inner FPPE.

Lemma 4.2. *In the metagame, for every pure Nash equilibrium, suppose p, x, t are the per-unit prices, allocation and payment of the induced FPPE. For every agent i, let $H_i \triangleq \{j \in [m] : x_{ij} > 0\}$ be the subset of items for which agent i receives a strictly positive fraction. Suppose agent i does not exhaust her true budget w_i, i.e., $t_i < w_i$, then*

$$\frac{S_i\left(\sum_{j\in[m]} \phi_{ij}x_{ij}\right)}{R_i\left(t_i\right)} \le \begin{cases} \frac{\sum_{j\in H_i} p_j}{\sum_{j\in H_i}(1-x_{ij})\phi_{ij}} & \text{if } H_i \ne \emptyset \\ \min_{j\in[m]} \frac{p_j}{\phi_{ij}} & \text{otherwise (i.e., } H_i = \emptyset) \end{cases}$$

where S_i and R_i are the derivative of valuation function V_i and money cost function C_i defined in Section 2, respectively.

One interpretation of Lemma 4.2 is as follows. First consider a hypothetical setting where item prices are fixed and the agent is acting as a price-taker, as in "standard" market equilibrium. Then it would be optimal for the agent to select items until the marginal price equals the marginal value-per-unit (i.e., $S(\cdot)$) divided by the marginal cost-per-unit (i.e., $R(\cdot)$), which can be formulated as the inequality in Lemma 4.2 with the discounting term $(1 - x_{ij})$ in the denominator of the right-hand side removed. However, in our setting, the agent is not acting as a price-taker: her behavior distorts the prices, and hence distorts the relationship between her allocation and the prices. Lemma 4.2 shows that this distortion is proportional to her allocation at equilibrium. If an agent is taking almost all of the items they care about (x_{ij} close to 1 for all j in H_i), the denominator on the right-hand side is close to 0 and the distortion is very large. In contrast, as long as a constant fraction of the items remains, the distortion is small. This is useful for our efficiency analysis: roughly speaking, an agent who gets a small allocation is acting approximately like a price-taker (and thus an approximate first-welfare-theorem analysis applies). On the other hand, an agent who gets a large allocation is anyway making a large contribution to the liquid welfare.

Now we are ready to prove Theorem 4.1.

Proof of Theorem 4.1. Fix an arbitrary pure Nash equilibrium and suppose p, x, t are the per-unit prices, allocation and payment of the inner FPPE, respectively. Let x^* be the optimal allocation that maximizes the liquid welfare. Consider the following partition $A_1 \bigsqcup A_2 \bigsqcup A_3$ of agents based on x, p, and x^*:

$$A_1 \triangleq \left\{i \in [n] : \sum_{j\in[m]} p_j x_{ij} = w_i\right\}$$

[17] Due to the symmetric of the instance, the same argument can be applied to $p_1 \le p_2$ and $p_2 > 0$ as well. The remaining case of $p_1 = p_2 = 0$ is already covered in the proof of Claim 3.7.

$$A_2 \triangleq \left\{i \in [n] : i \notin A_1 \wedge \sum\nolimits_{j\in[m]} \phi_{ij}x^*_{ij} \leq \sum\nolimits_{j\in[m]} \phi_{ij}x_{ij}\right\}$$
$$A_3 \triangleq \left\{i \in [n] : i \notin A_1 \wedge \sum\nolimits_{j\in[m]} \phi_{ij}x^*_{ij} > \sum\nolimits_{j\in[m]} \phi_{ij}x_{ij}\right\}$$

In words, A_1 corresponds to every agent i who exhausts her true budget w_i in the equilibrium; and A_2, A_3 correspond to the remaining agents divided according to whether their individual allocation is larger in the equilibrium outcome or in the optimal allocation.

In the following, we compare the willingness to pay (liquid welfare contribution) for agents from A_1, A_2, A_3 separately.

For every agent $i \in A_1$, note that $W_i(x^*_i) \leq w_i = W_i(x_i)$ where the inequality holds due to the definition of W_i, and the equality holds due to the definition of A_1.

For every agent $i \in A_2$, note that $W_i(x^*_i) \leq W_i(x_i)$ due to the fact that W_i is increasing and $\sum_{j\in[m]} \phi_{ij}x^*_{ij} \leq \sum_{j\in[m]} \phi_{ij}x_{ij}$ in the definition of A_2.

For every agent $i \in A_3$, let $H_i \triangleq \{j \in [m] : x_{ij} > 0\}$ be the subset of items for which agent i receives a strictly positive fraction. Consider two cases. First, suppose $H_i = \emptyset$. Note that

$$\frac{V_i\left(\sum_{j\in[m]} \phi_{ij}x^*_{ij}\right)}{C_i\left(\sum_{j\in[m]} p_jx^*_{ij}\right)} \overset{(a)}{\leq} \frac{\left(\sum_{j\in[m]} \phi_{ij}x^*_{ij}\right)\cdot S_i(0)}{\left(\sum_{j\in[m]} p_jx^*_{ij}\right)\cdot R_i(0)} \overset{(b)}{\leq} \frac{\left(\sum_{j\in[m]} \phi_{ij}x^*_{ij}\right)}{\left(\sum_{j\in[m]} p_jx^*_{ij}\right)}\cdot\left(\min_{j\in[m]}\frac{p_j}{\phi_{ij}}\right) \leq 1$$

where inequality (a) holds due to the concavity (convexity) of valuation function V_i (money cost function C_i); and inequality (b) holds due to Lemma 4.2. The above inequality further implies

$$W_i(x^*_i) = C_i^{-1}\left(V_i\left(\sum_{j\in[m]}\phi_{ij}x^*_{ij}\right)\right) \leq \sum_{j\in[m]} p_jx^*_{ij}$$

Next, suppose $H_i \neq \emptyset$. Let $\Delta_i \triangleq \frac{\sum_{j\in H_i} p_j}{\sum_{j\in H_i}\phi_{ij}}$. Due to the "highest bang-per-buck" property of FPPE, $\frac{p_j}{\phi_{ij}} \geq \Delta_i$ for every item $j \in [m]$, and equality holds for $j \in H_i$. Note that

$$\begin{aligned}
W_i(x^*_i) &\overset{(a)}{\leq} W_i(x_i) + \frac{V_i\left(\sum_{j\in[m]}\phi_{ij}x^*_{ij}\right) - V_i\left(\sum_{j\in[m]}\phi_{ij}x_{ij}\right)}{R_i(W_i(x_i))}\\
&\overset{(b)}{\leq} W_i(x_i) + \frac{V_i\left(\sum_{j\in[m]}\phi_{ij}x^*_{ij}\right) - V_i\left(\sum_{j\in[m]}\phi_{ij}x_{ij}\right)}{R_i(t_i)}\\
&\overset{(c)}{\leq} W_i(x_i) + \frac{S_i\left(\sum_{j\in[m]}\phi_{ij}x_{ij}\right)}{R_i(t_i)}\cdot\sum\nolimits_{j\in[m]}\phi_{ij}(x^*_{ij} - x_{ij})\\
&\overset{(d)}{\leq} W_i(x_i) + \frac{\sum_{j\in H_i} p_j}{\sum_{j\in H_i}(1 - x_{ij})\phi_{ij}}\cdot\sum\nolimits_{j\in[m]}\phi_{ij}(x^*_{ij} - x_{ij})\\
&\overset{(e)}{\leq} W_i(x_i) + \frac{\sum_{j\in H_i} p_j}{\sum_{j\in H_i}\phi_{ij}}\cdot\sum\nolimits_{j\in[m]}\phi_{ij}x^*_{ij}\\
&\overset{(f)}{=} W_i(x_i) + \Delta_i\cdot\sum\nolimits_{j\in[m]}\phi_{ij}x^*_{ij}\\
&\overset{(g)}{\leq} W_i(x_i) + \sum\nolimits_{j\in[m]} p_jx^*_{ij}
\end{aligned}$$

where inequality (a) holds due to the definition of W_i and the convexity of money cost function C_i; inequality (b) holds since R_i is weakly increasing implied by the convexity of C_i, and $W_i(x_i) \geq t_i$ which is implied by the definition of W_i and the fact that agent i's utility is non-negative in the equilibrium; inequality (c) holds due to the concavity of valuation function V_i; inequality (d) holds due to Lemma 4.2; and inequality (e) holds since $\frac{\sum_{j\in[m]}\phi_{ij}(x^*_{ij}-x_{ij})}{\sum_{j\in[m]}\phi_{ij}(1-x_{ij})} \leq \frac{\sum_{j\in[m]}\phi_{ij}x^*_{ij}}{\sum_{j\in[m]}\phi_{ij}}$ by algebra; equality (f) holds due to the definition of Δ_i; and inequality (g) holds since $\frac{p_j}{\phi_{ij}} \geq \Delta_i$ for every item $j \in [m]$.

Putting pieces together, we have the following upper bound of the optimal liquid welfare $W(x^*)$:

$$\begin{aligned}
W(x^*) &= \sum\nolimits_{i\in[n]} W_i(x^*_i)\\
&\leq \sum\nolimits_{i\in A_3}\sum\nolimits_{j\in[m]} p_jx^*_{ij} + \sum\nolimits_{i\in A_1\sqcup A_2\sqcup A_3} W_i(x_i)\\
&\overset{(a)}{\leq} \sum\nolimits_{j\in[m]} p_j + \sum\nolimits_{i\in A_1\sqcup A_2\sqcup A_3} W_i(x_i) \overset{(b)}{\leq} 2W(\mathbf{x})
\end{aligned}$$

where inequality (a) holds since $\sum_{i\in A_3} x^*_{ij} \leq 1$ for every item $j \in [m]$; and inequality (b) holds since $\sum_{j\in[m]} p_j \leq W(\mathbf{x})$ since all agents receive non-negative utility in the equilibrium. □

4.2 Mixed Nash Equilibrium

Now we present the PoA for the mixed equilibria of the metagame.

Theorem 4.3. *In the metagame, the price of anarchy $\Gamma_{\texttt{Mixed}}$ under mixed Nash equilibrium is in [2, 4].*

Note that Example 4.1 also serves as a lower bound for Theorem 4.3, so what remains is to prove the upper bound. We first provide some intuition into the high-level approach. In our analysis of pure Nash equilibrium, we analyzed the impact of small budget adjustments given the messages of the other agents. However, in a mixed Nash equilibrium, the messages (and hence outcomes and prices) may be random, so the impact of local adjustments less clear. We instead consider a specific budget-setting strategy that each agent will consider as a deviation. Namely, each agent considers the *expected* per-unit prices of the inner FPPE if she were not present, and then sets a budget equal to the expected total payment of her part of the optimal allocation under those prices. We use Lemma 2.2 and some other properties of FPPE to characterize the per-unit prices of the inner FPPE and the agent's utility under such a deviation. Consequently, we obtain an upper bound on the optimal liquid welfare.

5 EXTENSION I: BAYESIAN ENVIRONMENTS

In this section, we explore the generalization of our model and results to from full information environments to Bayesian environments. In this *Bayesian metagame* extension, rather than assuming that each agent i has a fixed type (V_i, w_i, C_i), we assume that each agent i's type is independently drawn from a type distribution F_i. In the following, we formally define the solution concept – Bayesian Nash equilibrium and price of anarchy under Bayesian Nash equilibrium. Finally, we prove that the price of anarchy under Bayesian Nash equilibrium is between $[2, 4]$ in Theorem 5.1.

In Bayesian metagame, a strategy $\mathbf{s}_i$ of agent i is a stochastic mapping from agent i's type (V_i, w_i, C_i) to a randomized message $(\tilde{\mathbf{v}}_i, \tilde{\mathbf{w}}_i)$. The *Bayesian Nash equilibrium* is defined as follows.

Definition 5.1 (Bayesian Nash equilibrium). *For agents with type distributions $\{F_i\}_{i\in[n]}$, a Bayesian Nash equilibrium is a strategy profile $\{\mathbf{s}_i\}_{i\in[n]}$ such that for every agent i, every realized type (V_i, w_i, C_i), and every message $(\tilde{v}_i^\dagger, \tilde{w}_i^\dagger)$,*

$$\mathbb{E}[u_i(\tilde{v}_i, \tilde{w}_i, \tilde{v}_{-i}, \tilde{w}_{-i})] \geq \mathbb{E}\Big[u_i(\tilde{v}_i^\dagger, \tilde{w}_i^\dagger, \tilde{v}_{-i}, \tilde{w}_{-i})\Big]$$

where the expectation is taken over agent i's strategy $(\tilde{v}_i, \tilde{w}_i) \sim \mathbf{s}_i(V_i, w_i, C_i)$, other agents' types $(V_{-i}, w_{-i}, C_{-i}) \sim F_{-i}$ and corresponding strategies $(\tilde{v}_{-i}, \tilde{w}_{-i}) \sim \mathbf{s}_{-i}(V_{-i}, w_{-i}, C_{-i})$.

Clearly, the Bayesian metagame is a generalization of our baseline model, as every problem instance in the baseline model where agents have fixed types can be viewed as a problem instance where agents' types are drawn from single point-mass distributions in the Bayesian metagame. As a result of this equivalence, Bayesian Nash equilibrium also encompasses mixed Nash equilibrium as a generalization.

Similar to the baseline model, we evaluate the performance of the Bayesian metagame by measuring the approximation of liquid welfare. This is done by comparing the worst expected liquid welfare among all possible equilibria with the optimal expected liquid welfare over the randomness of agents' type and messages, and taking the supremum over all instances.

Definition 5.2 (Price of anarchy in Bayesian environments). *The* price of anarchy (PoA) *of the metagame* $\Gamma_{\texttt{Bayes}}$ *under Bayesian Nash equilibrium is*

$$\Gamma_{\texttt{Bayes}} \triangleq \sup_{n,m,\phi} \sup_{\{F_i\}_{i\in[n]}} \frac{\mathbb{E}\big[\max_x W(x \mid \{(V_i, w_i, C_i)\}_{i\in[n]})\big]}{\inf_{\mathbf{s}\in\texttt{Bayes}} \hat{W}(\mathbf{s} \mid \{F_i\}_{i\in[n]})}$$

Here the expectation in the numerator is taken over random type $\{(V_i, w_i, C_i)\}_{i\in[n]} \sim \{F_i\}_{i\in[n]}$, and $W(x \mid \{(V_i, w_i, C_i)\}_{i\in[n]})$ is the liquid welfare of allocation x given agents' types $\{(V_i, w_i, C_i)\}_{i\in[n]}$. In the denominator, $\texttt{Bayes}$ *is the set of strategy profiles in all Bayesian Nash equilibrium given type distributions $\{F\}_{i\in[n]}$, and $\hat{W}(\mathbf{s} \mid \{F_i\}_{i\in[n]})$ is the expected liquid welfare under strategy profile* $\mathbf{s}$ *defined as*

$$\hat{W}(\mathbf{s} \mid \{F_i\}_{i\in[n]}) \triangleq \sum_{i\in[n]} \hat{W}_i(\mathbf{s} \mid \{F_i\}_{i\in[n]}),$$

$$\hat{W}_i(\mathbf{s} \mid \{F_\ell\}_{\ell\in[n]}) \triangleq \mathbb{E}_{(V_i,w_i,C_i)\sim F_i}[W_i(\mathbf{x}_i(\mathbf{s}, (V_i, w_i, C_i), F_{-i}))]$$

for every agent $i \in [n]$. Here $\mathbf{x}_i(\mathbf{s}, (V_i, w_i, C_i), F_{-i})$ is the randomized allocation of agent i with realized type (V_i, w_i, C_i) when agents report under strategy $\mathbf{s}$*, and the randomness is over agent i's message, other agents' types and their messages.*

The main result of this section is the bound on the PoA under Bayesian Nash equilibrium.

Theorem 5.1. *In the Bayesian metagame, the PoA under Bayesian Nash equilibrium lies in* $[2, 4]$.

6 EXTENSION II: RESTRICTED MESSAGES

In this section, our goal is to address the following question: *In the metagame, what information encoded in agents' messages is essential and cannot be disregarded?* Answering this question holds significant practical implications. For instance, in digital advertising markets, the platform (seller) possesses control over the design of the auction interface for advertisers (agents). As a reminder, our metagame draws inspiration from such an interface where advertisers declare both budget and maximum bid constraints to their autobidder.

In the proof of Lemma 3.2, we establish that every message $(\tilde{v}_i, \tilde{w}_i)$ can be dominated by a message $(\tilde{v}_i^\dagger, \tilde{w}_i^\dagger)$ with reported value $\tilde{v}_i^\dagger = \infty$. Similarly, in the proofs of other results in previous sections, we often construct deviation strategy with reported value $\tilde{v}_i^\dagger = \infty$ and carefully design reported budget $\tilde{w}_i^\dagger$. Loosely speaking, this indicates that for agents, *the strategic decision of their reported budget $\tilde{w}$ holds greater importance than their reported value (aka maximum bid) $\tilde{v}$.*

Motivated by this intuition, we proceed to investigate two variants of the metagame, wherein agents report either only budgets or only values. In Section 6.1, we extend our analysis of the price of anarchy to the variant metagame where agents only report budgets. On the other hand, in Section 6.2, we present a negative result demonstrating that the price of anarchy under Bayesian Nash equilibrium can be as high as linear in the number of agents in the variant metagame where agents only report values.

6.1 Metagame with Budget Reporting Only

In this subsection, we examine a variant of the metagame where agents exclusively report their budgets. The definitions of pure Nash equilibrium, mixed Nash equilibrium, and Bayesian Nash equilibrium are adjusted accordingly to accommodate this variant.

Definition 6.1 (Metagame with budget reporting only). *Each agent i decides on a message $\tilde{w}_i \in \mathbb{R}_+^\infty$ reported to the seller. Given reported message profile $\{\tilde{w}_i\}_{i\in[n]}$, the seller implements allocation $x(\{\tilde{w}_i\}_{i\in[n]})$ and payment $t(\{\tilde{w}_i\}_{i\in[n]})$ induced by the FPPE, assuming that agents have budgeted utility with types $\{(\infty, \tilde{w}_i)\}_{i\in[n]}$.*

In words, in this variant of the metagame with budget reporting only, the seller treats each agent i with general utility model as a budgeted agent with value per click $v_i = \infty$ and budget $w_i = \tilde{w}_i$ reported from the agent. In this variant, we obtain the same price of anarchy (PoA) bounds as the original metagame where agents report both values and budgets.

Proposition 6.1. *In the metagame with budget reporting only, the price of anarchy* $\Gamma_{\texttt{Pure}}$ *under pure Nash equilibrium is 2, and the price of anarchy* $\Gamma_{\texttt{Mixed}}$ ($\Gamma_{\texttt{Bayes}}$) *under mixed (Bayesian) Nash equilibrium is between [2, 4].*

Another variant for agents with linear valuation functions. In the remaining of this subsection, we restrict out attention to agents with linear valuation functions, i.e., $V_i(\sum_j \phi_{ij}x_{ij}) = v_i \cdot \sum_j \phi_{ij}x_{ij}$ where v_i is her *value per click*. However, agents may still have hard budget w_i and differentiable, weakly increasing, weakly convex money cost function C_i. For such agents, we denote (v_i, w_i, C_i) by their types. We consider another variant of the metagame with budget reporting only as follows.

Definition 6.2 (Metagame with budget reporting only and known linear valuations). *Each agent i with type (v_i, w_i, C_i) decides on a message $\tilde{w}_i \in \mathbb{R}_+^\infty$ reported to the seller. Given reported message profile $\{\tilde{w}_i\}_{i\in[n]}$, the seller implements allocation $x(\{\tilde{w}_i\}_{i\in[n]})$ and payment $t(\{\tilde{w}_i\}_{i\in[n]})$ induced by the FPPE, assuming that agents have budgeted utility with types $\{(v_i, \tilde{w}_i)\}_{i\in[n]}$.*

In words, in this variant of the metagame with budget reporting only, the seller knows the value per click v_i of each agent i and treats this agent with type (v_i, w_i, C_i) as a budgeted agent with value per click v_i and budget $w_i = \tilde{w}_i$ reported from the agent. In this variant, we still obtain the same price of anarchy (PoA) bounds as the original metagame where agents report both values and budgets.

Proposition 6.2. *In the metagame with budget reporting only and known linear valuations, the price of anarchy $\Gamma_{\texttt{Pure}}$ under pure Nash equilibrium is 2, and the price of anarchy $\Gamma_{\texttt{Mixed}}$ ($\Gamma_{\texttt{Bayes}}$) under mixed (Bayesian) Nash equilibrium is between [2, 4].*

6.2 Metagame with Value Reporting Only

Let us examine a variant of the metagame where agents only report their values (i.e., the maximum bid).

Definition 6.3 (Metagame with value reporting only). *Each agent i decides on a message $\tilde{v}_i \in \mathbb{R}_+^\infty$ reported to the seller. Given reported message profile $\{\tilde{v}_i\}_{i\in[n]}$, the seller implements allocation $x(\{\tilde{v}_i\}_{i\in[n]})$ and payment $t(\{\tilde{v}_i\}_{i\in[n]})$ induced by the FPPE, assuming that agents have budgeted utility with types $\{(\tilde{v}_i, \infty)\}_{i\in[n]}$.*

In words, in this variant of the metagame with value reporting only, the seller treats each agent i with general utility model as a budgeted agent with value per click $v_i = \tilde{v}_i$ reported from the agent and budget $w_i = \infty$ reported from the agent. In this variant, we present the following negative result on the price of anarchy under Bayesian Nash equilibrium.

Proposition 6.3. *In the metagame with value reporting only, the price of anarchy $\Gamma_{\texttt{Bayes}}$ under Bayesian Nash equilibrium is at least $\Omega(n)$, even for budgeted agents and a single item.*

7 SINGLE-ITEM INSTANCES

To shed light on agents' strategic behavior in the metagame, this section focuses on single-item instances. We characterize the pure Nash equilibrium for single-item instances. All results and analysis hold for more general instances with homogeneous items: each agent i has the same click-through rate ϕ_{ij} for all item j, i.e., $\phi_{ij} = \phi_{ij'}$ for every $j, j' \in [m]$, $i \in [n]$.

By restricting to single-item instances, the remaining of this section drops subscript index j for the item. Moreover, without loss of generality, we assume $\boldsymbol{\phi}_i = 1$ for all agents and drop it as well.

Before presenting the main result of this section, we introduce the following two auxiliary notations that are used in the equilibrium characterization (Theorem 7.2): fix a per-unit price $p \in [0, \infty)$, define

$$y_i(p) \triangleq \frac{w_i}{p} \wedge \min\{x_i \in [0,1] : S_i(x_i) \cdot (1 - x_i) \le p \cdot R_i(px_i)\}$$

$$z_i(p) \triangleq \frac{w_i}{p} \wedge \max\{x_i \in [0,1] : S_i(x_i) \ge p \cdot R_i(px_i)\}$$

Loosely speaking, $y_i(p)$ is the smallest allocation such that the agent i has no incentive to weakly increase her reported budget in the metagame with induced per-unit price p. Specifically, the inequality in the definition $y_i(p)$ is exactly the inequality of Lemma 4.2 for single-item instances. On the other side, $z_i(p)$ is the largest allocation such that the agent i has a weakly positive marginal utility when facing a fixed per-unit price p. If $S_i(x_i) < p \cdot R_i(px_i)$ for all $x_i \in [0, 1]$, we set $z_i(p) = 0$. Since valuation function V_i (money cost function C_i) is differentiable, weakly increasing, weakly concave (convex) and $V_i(0) = 0$ ($C_i(0) = 0$), we make the following observation about $y_i(p)$ and $z_i(p)$.

Observation 7.1. *For every agent $i \in [n]$, functions $y_i(p)$ and $z_i(p)$ satisfy the following properties:*

1. *Both $y_i(p)$ and $z_i(p)$ are continuous and weakly decreasing in p.*
2. *Both $y_i(0) = z_i(0) = 1$, $\lim_{p\to\infty} y_i(p) = \lim_{p\to\infty} z_i(p) = 0$, and $y_i(x_i) \le z_i(x_i)$ for all $x_i \in \mathbb{R}_+$.*

We now present the main result of the section.

Theorem 7.2. *In the metagame with a single item, pure Nash equilibrium always exists. Specifically, there exist two types of equilibrium:*

1. *(Low-price equilibrium) Define non-empty subinterval P_L as*
$$P_L \triangleq \left\{p \in [0,\infty) : \sum\nolimits_{i\in[n]} y_i(p) = 1\right\}$$
For every $p \in P_L$, there exists pure Nash equilibrium whose inner FPPE has per-unit price p. Such equilibrium can be induced by reported message profile $\{(\tilde{v}_i, \tilde{w}_i)\}_{i\in[n]}$ constructed as
$$\forall i \in [n]: \qquad \tilde{v}_i \triangleq \infty,\ \tilde{w}_i \triangleq p y_i(p)$$
Moreover, $\min P_L$ is the lowest per-unit price in all pure Nash equilibrium.
2. *(High-price equilibrium) Define non-empty subinterval P_H as*
$$P_H \triangleq \left\{p \in [0,\infty) : \sum\nolimits_{i\in[n]} y_i(p) \le 1 \text{ and } \sum\nolimits_{i\in[n]} z_i(p) \ge 1\right\}$$
For every $p \in P_H$, if there exists agent i^* such that $y_{i^*}(p) + \sum_{i\in[n]:i\ne i^*} z_i(p) \ge 1$, then there exists pure Nash equilibrium whose inner FPPE has per-unit price p. Such equilibrium can be induced by reported message profile $\{(\tilde{v}_i, \tilde{w}_i)\}_{i\in[n]}$ constructed as
$$i = i^*: \qquad \tilde{v}_i \triangleq p,\ \tilde{w}_i \triangleq \infty,$$
$$\forall i \ne i^*: \qquad \tilde{v}_i \triangleq \infty,\ \tilde{w}_i \triangleq p\hat{x}_i$$
where $\{\hat{x}_i\}_{i\ne i^*}$ is an arbitrary solution such that $\sum_{i\ne i^*} \hat{x}_i = 1 - y_{i^*}(p)$ and $y_i(p) \le \hat{x}_i \le z_i(p)$. Moreover, $\max P_H$ is the highest possible per-unit price in all pure Nash equilibrium.

By analyzing a two-item budgeted-agent instance from Example 2.3, Section 3.2 shows the non-existence of pure Nash equilibrium for general instances. In contrast, Theorem 7.2 confirms the existence of pure Nash equilibrium for single-item instances.

Since $y_i(p) \le z_i(p)$ for all $p \in \mathbb{R}_+$ (Observation 7.1), it is straightforward to verify that $P_L \subseteq P_H$. Moreover, for every $p \in P_L$, it can be constructed as both low-price equilibrium and high-price equilibrium (since condition "$\exists i^*, y_{i^*}(p) + \sum_{i\ne i^*} z_i(p) \ge 1$" is satisfied trivially). A natural question is whether there exists high-price equilibrium with per-unit price $p \in P_H$ such that there exists no low-price equilibrium with the same per-unit price, i.e., $p \in P_L$. The answer is yes. Consider Example 2.2:

$$y_i(p) = \max\left\{1 - \frac{p}{v_i}, 0\right\}, \qquad z_i(p) = \mathbb{I}\{p \le v_i\}$$
$$P_L = \left\{\frac{v_1 v_2}{v_1+v_2}\right\}, \qquad P_H = \left[\frac{v_1 v_2}{v_1+v_2}, v_1\right].$$

Though the reported messages are slightly different, equilibrium with per-unit price v_2 (resp. $\frac{v_1 v_2}{v_1+v_2}$) described in Claim 3.4 (resp. Claim 3.5) is equivalent to a high-price (resp. low-price) equilibrium. In fact, we can extend Example 2.2 and construct natural scenario with multiple pure Nash equilibria. In those equilibria, the per-unit prices of inner FPPE are different. Consequently, a fixed agent's utility is different in different equilibrium.

Proposition 7.3. *In the metagame, for budgeted (or linear) agents with type $\{(v_i, w_i)\}$, if $w_i > \frac{1}{4}v_i$ for all agents, then there exists multiple pure Nash equilibrium. Specifically, besides low-price equilibrium and high-price equilibrium with per-unit price in P_L, there also exists high-price equilibrium with per-unit price in $P_H \backslash P_L$.*

By utilizing the monotonicity of $y_i(\cdot), z_i(\cdot)$, we can compute P_L, P_H and thus construct equilibrium in Theorem 7.2 in polynomial time. See the full version for more details.

REFERENCES

[1] Gagan Aggarwal, Ashwinkumar Badanidiyuru, and Aranyak Mehta. 2019. Autobidding with constraints. In *Web and Internet Economics: 15th International Conference, WINE 2019, New York, NY, USA, December 10–12, 2019, Proceedings 15*. Springer, 17–30.
[2] Airbnb. 2023. *Smart Pricing*.
[3] Yeganeh Alimohammadi, Aranyak Mehta, and Andres Perlroth. 2023. Incentive Compatibility in the Auto-bidding World. *Proceedings of the 24th ACM Conference on Economics and Computation* (2023).
[4] Amazon. 2023. *Adjust pricing quickly and automatically*.
[5] Yossi Azar, Michal Feldman, Nick Gravin, and Alan Roytman. 2017. Liquid price of anarchy. In *Algorithmic Game Theory: 10th International Symposium, SAGT 2017, L'Aquila, Italy, September 12–14, 2017, Proceedings*. Springer, 3–15.
[6] Moshe Babaioff, Richard Cole, Jason Hartline, Nicole Immorlica, and Brendan Lucier. 2021. Non-Quasi-Linear Agents in Quasi-Linear Mechanisms. In *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)*, Vol. 185. 84.
[7] Santiago Balseiro, Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. 2021. Robust auction design in the auto-bidding world. *Advances in Neural Information Processing Systems* 34 (2021), 17777–17788.
[8] Santiago Balseiro, Anthony Kim, Mohammad Mahdian, and Vahab Mirrokni. 2017. Budget management strategies in repeated auctions. In *Proceedings of the 26th International Conference on World Wide Web*. 15–23.
[9] Santiago Balseiro, Haihao Lu, and Vahab Mirrokni. 2020. Dual mirror descent for online allocation problems. In *International Conference on Machine Learning*. PMLR, 613–628.
[10] Santiago R Balseiro, Yuan Deng, Jieming Mao, Vahab S Mirrokni, and Song Zuo. 2021. The landscape of auto-bidding auctions: Value versus utility maximization. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 132–133.
[11] Santiago R Balseiro and Yonatan Gur. 2019. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science* 65, 9 (2019), 3952–3968.
[12] Santiago R Balseiro, Haihao Lu, and Vahab Mirrokni. 2023. The best of many worlds: Dual mirror descent for online allocation problems. *Operations Research* 71, 1 (2023), 101–119.
[13] C. Borgs, J. Chayes, N. Immorlica, M. Mahdian, and A. Saberi. 2005. Multi-unit auctions with budget-constrained bidders. In *Proc. 6th ACM Conf. on Electronic Commerce*.
[14] Christian Borgs, Jennifer T. Chayes, Nicole Immorlica, Kamal Jain, Omid Etesami, and Mohammad Mahdian. 2007. Dynamics of bid optimization in online advertisement auctions. In *16thInternational World Wide Web Conference*. 531–540.
[15] Ioannis Caragiannis and Alexandros A Voudouris. 2016. Welfare guarantees for proportional allocations. *Theory of Computing Systems* 59 (2016), 581–599.
[16] Ioannis Caragiannis and Alexandros A Voudouris. 2018. The efficiency of resource allocation mechanisms for budget-constrained users. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. 681–698.
[17] Matteo Castiglioni, Andrea Celli, Alberto Marchesi, Giulia Romano, and Nicola Gatti. 2022. A unifying framework for online optimization with long-term constraints. *Advances in Neural Information Processing Systems* 35 (2022), 33589–33602.
[18] Xi Chen, Christian Kroer, and Rachitesh Kumar. 2021. The Complexity of Pacing for Second-Price Auctions. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 318–318.
[19] George Christodoulou, Alkmini Sgouritsa, and Bo Tang. 2015. On the efficiency of the proportional allocation mechanism for divisible resources. In *Algorithmic Game Theory: 8th International Symposium, SAGT 2015, Saarbrücken, Germany, September 28-30, 2015. Proceedings 8*. Springer, 165–177.
[20] Vincent Conitzer, Christian Kroer, Debmalya Panigrahi, Okke Schrijvers, Nicolas E Stier-Moses, Eric Sodomka, and Christopher A Wilkens. 2022. Pacing Equilibrium in First Price Auction Markets. *Management Science* (2022).
[21] Vincent Conitzer, Christian Kroer, Eric Sodomka, and Nicolas E Stier-Moses. 2022. Multiplicative pacing equilibria in auction markets. *Operations Research* 70, 2 (2022), 963–989.
[22] Yuan Deng, Mohammad Mahdian, Jieming Mao, Vahab Mirrokni, Hanrui Zhang, and Song Zuo. 2023. Efficiency of the Generalized Second-Price Auction for Value Maximizers. *arXiv preprint arXiv:2310.03105* (2023).
[23] Yuan Deng, Jieming Mao, Vahab Mirrokni, Hanrui Zhang, and Song Zuo. 2022. Efficiency of the first-price auction in the autobidding world. *arXiv preprint arXiv:2208.10650* (2022).
[24] Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. 2021. Towards efficient auctions in an auto-bidding world. In *Proceedings of the Web Conference 2021*. 3965–3973.
[25] Shahar Dobzinski and Renato Paes Leme. 2014. Efficiency guarantees in auctions with budgets. In *International Colloquium on Automata, Languages, and Programming*. Springer, 392–404.
[26] Expedia. 2023. *Find our lowest price on your next flight with Price Tracking*.
[27] Giannis Fikioris and Éva Tardos. 2023. Liquid welfare guarantees for no-regret learning in sequential budgeted auctions. In *Proceedings of the 24th ACM Conference on Economics and Computation*. 678–698.
[28] Jason Gaitonde, Yingkai Li, Bar Light, Brendan Lucier, and Aleksandrs Slivkins. 2023. Budget Pacing in Repeated Auctions: Regret and Efficiency Without Convergence. In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, Vol. 251. 52.
[29] Negin Golrezaei, Patrick Jaillet, Jason Cheuk Nam Liang, and Vahab Mirrokni. 2023. Pricing against a Budget and ROI Constrained Buyer. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 9282–9307.
[30] Google. 2023. *About automated bidding*.
[31] John Richard Hicks et al. 1975. Value and capital: An inquiry into some fundamental principles of economic theory. *OUP Catalogue* (1975).
[32] Hopper. 2023. *How do Price Predictions work with Hopper?*
[33] Yoav Kolumbus and Noam Nisan. 2022. Auctions between regret-minimizing agents. In *Proceedings of the ACM Web Conference 2022*. 100–111.
[34] Yoav Kolumbus and Noam Nisan. 2022. How and why to manipulate your own agent: On the incentives of users of learning agents. *Advances in Neural Information Processing Systems* 35 (2022), 28080–28094.
[35] Juncheng Li and Pingzhong Tang. 2022. Auto-bidding Equilibrium in ROI-Constrained Online Advertising Markets. *arXiv preprint arXiv:2210.06107* (2022).
[36] Christopher Liaw, Aranyak Mehta, and Andres Perlroth. 2022. Efficiency of non-truthful auctions under auto-bidding. *arXiv preprint arXiv:2207.03630* (2022).
[37] Pinyan Lu and Tao Xiao. 2015. Improved efficiency guarantees in auctions with budgets. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. 397–413.
[38] Brendan Lucier, Sarath Pattathil, Aleksandrs Slivkins, and Mengxiao Zhang. 2023. Autobidders with budget and roi constraints: Efficiency, regret, and pacing dynamics. *arXiv preprint arXiv:2301.13306* (2023).
[39] Aranyak Mehta. 2022. Auction design in an auto-bidding setting: Randomization improves efficiency beyond vcg. In *Proceedings of the ACM Web Conference 2022*. 173–181.
[40] Aranyak Mehta and Andres Perlroth. 2023. Auctions without commitment in the auto-bidding world. *arXiv preprint arXiv:2301.07312* (2023).
[41] Meta. 2023. *About Meta Bid Strategies*.
[42] Microsoft. 2023. *Automated bidding*.
[43] W. Vickrey. 1961. Counterspeculation, Auctions, and Competitive Sealed Tenders. *J. of Finance* 16 (1961), 8–37.